


Doc#: 395003.1
RDJ/GHF/ecf
May 23, 2003

PATENT APPLICATION
Attorney's Docket No.: 1958.1018-000

[Handwritten signature]
#18
7/29/03



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: Richard L. Frank, Michael J. Cusson, Joydip Kundu, and Daniel E. O'Shaughnessy

Application No.: 09/321,090 Group: 2153

Filed: May 28, 1999 Examiner: Walter Benson

For: A QUORUMLESS CLUSTER USING DISK-BASED MESSAGING

CERTIFICATE OF MAILING	
I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as First Class Mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450	
on <u>May 28, 2003</u>	<u>Elaine Leahy</u>
Date	Signature
<u>Elaine Leahy</u>	
Typed or printed name of person signing certificate	

DECLARATION OF RICHARD L. FRANK
UNDER 37 C.F.R. §1.131

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

RECEIVED
JUN -5 2003
TECHNOLOGY CENTER 2800

Sir:
I, Richard L. Frank, of 25 Brown Lane, Groton, Massachusetts 01450 declare, and state that:

1. I am a co-inventor of the above-captioned application, filed on May 28, 1999, and assigned to Oracle Corporation.

OID-1999-035-03

714.03(a)
II

2. The inventors of the above-captioned application are Michael J. Cusson, Joydip Kundu, Daniel E. O'Shaughnessy and myself.
3. I have personal knowledge surrounding the conception and reduction to practice of the claimed invention.
4. The idea of granting membership in a network cluster to a node if the node has access to the shareable storage device, described in Claims 1-21, was conceived and reduced to practice prior to January 19, 1999.
5. Prior to January 19, 1999, software code that implemented the invention of Claims 1-21 was written and tested. As evidence of this reduction to practice, Exhibits A-E are enclosed herewith. Exhibits A-E are true copies of selected pages of a version of the software code, obtained from Oracle's software management systems and dated before January 19, 1999. In accordance with accepted U.S. Patent Office practice, portions of the source code in Exhibits A-E have been redacted.
6. In Exhibits A-E, several routines are presented, which are used to grant membership in the network cluster. Examples of these routines include InitCm (Exhibit A) CreateCmThread, (Exhibit B) CmStartUp, (Exhibit C), InitControlDisk (Exhibit D) and WaitForCM (Exhibit E).

By way of background, each node executes cluster manager (CM) software, which it uses to establish access to the shareable storage device and membership in the network cluster. In the particular implementation described here, the shareable storage device is implemented as a bound set of disk drives. A node requesting membership in the network cluster executes its cluster manager, which calls the routine InitCm (Exhibit A). InitCm initializes the cluster

manager and spawns a cluster manager thread by running the routine CreateCmThread (Exhibit B).

The routine CreateCMThread calls CmStartUp (Exhibit C), which reads the cluster definition and initializes the disk-based messaging area (control area) on the shareable storage device (control disk). CmStartUp invokes the routine InitControlDisk (Exhibit D) to determine, among other things, the valid members of the cluster. To be a valid member, a node must be able to access the shareable storage device.

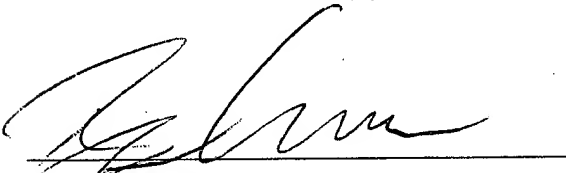
For example, InitControlDisk reads the shareable storage device to determine if any nodes have written to their message location. In particular, InitControlDisk uses the routine CmReadAnyFromBoundSet to determine which nodes have written to their message location on the shareable storage device. If the node requesting membership in the cluster cannot access the shareable storage device, then it cannot be a member of the cluster (i.e., status != NEDC_IO_SUCCESS). Thus, when the "status!= NEDC_IO_SUCCESS" statement is processed by the CmStartUp routine, CmStartUp provides a pointer in the node's local registry (e.g. pClub) to CmAbort. This indicates that the node's attempt to join the cluster has failed, and that the node's cluster manager processes are to be shut down.

If the node, however, can access the shareable storage device, it can be granted membership in the cluster. In this situation, the node satisfies the condition that its "status!= NEDC_IO_SUCCESS," and that its "status=WriteControlArea," which indicates that the node is able to access the shareable storage device. This status is processed by the InitCm routine to confirm that the creation of the cluster manager thread was successful. The routine WaitForCm (Exhibit E) is then called by InitCm to confirm that the cluster manager is initialized (e.g. return true). When InitCm receives this indication, the node's request for membership in the cluster is accepted (e.g. ready to accept application registration).

Note that the node's membership in the network cluster is predicated on its ability to access the shareable storage and not its network connectivity to other nodes. As a result, the cluster can be formed from a single node. Also, unlike prior cluster management systems, a quorum of nodes is not required to form a cluster.

7. The software code referenced in Sections 5 and 6 of this Declaration shows that prior to January 19, 1999, we implemented a cluster membership management technique in which membership is granted based on access to the shareable storage device.

I declare further that all statements made herein of my own knowledge are true and that all statements made on information and belief are true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.


Richard L. Frank

27 May 2003
Date